

CLAIMS

What is claimed is:

1. A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:
 - partitioning a file F into a plurality of subfiles;
 - performing distribution of said file F to a plurality of recipient nodes using a distribution technique that comprises
 - (a) attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes, wherein the first node attempts to communicate at least one subfile to each recipient node of said first group but not all of said plurality of subfiles to any recipient node of said first group, and
 - (b) said plurality of recipient nodes of said first group attempting to exchange their respective subfiles received from said first node, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;
 - detecting a failed node of said plurality of recipient nodes; and
 - said distribution technique adapting to distribute all of the subfiles of said file F to each non-failed node of said plurality of recipient nodes.
2. The method of claim 1 wherein said distribution technique adapting responsive to said detecting a failed node.
3. The method of claim 1 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:
 - attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group.
4. The method of claim 1 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:
 - attempting to distribute the plurality of subfiles from said first node to said plurality of recipient nodes of said first group concurrently.

5. The method of claim 1 wherein said plurality of recipient nodes of said first group attempting to exchange their respective subfiles further comprises:

each of said plurality of recipient nodes attempting to establishing concurrent communication connections to every other recipient node of said first group.

6. The method of claim 1 wherein said detecting a failed node comprises said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node.

7. The method of claim 6 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises said first node attempting to establish concurrent communication connections to the recipient nodes of said first group, and wherein said distribution technique adapting comprises:

responsive to said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node, said first node using its established concurrent communication connections with non-failed nodes of said first group to communicate the particular subfile to said non-failed nodes.

8. The method of claim 6 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises said first node attempting to establish concurrent communication connections to the recipient nodes of said first group, and wherein said distribution technique adapting comprises:

responsive to said first node detecting a failed node in said first group such that said first node is unable to communicate a particular subfile to such failed node, said first node triggering a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate the particular subfile to said non-failed nodes.

9. The method of claim 1 wherein said detecting a failed node comprises:
said recipient nodes of said first group exchanging heartbeat messages;
at least one recipient node of said first group detecting a failed node from analysis of heartbeat messages received; and

said at least one recipient node of said first group notifying said first node of said detected failed node.

10. The method of claim 1 wherein said detecting a failed node comprises:
said non-failed recipient nodes of said first group sending heartbeat messages to said first node; and

said first node detecting a failed node from analysis of received heartbeat messages from said non-failed recipient nodes.

11. The method of claim 1 further comprising:
said first group of recipient nodes attempting to communicate said file *F* to a second group comprising a plurality of recipient nodes.

12. The method of claim 11 further comprising:
each recipient node of said first group attempting to communicate a subfile to at least one recipient node of said second group.

13. The method of claim 12 further comprising:
each recipient node of said first group attempting to communicate the subfile that it received from said first node to a corresponding node of the second group.

14. The method of claim 12 wherein said detecting a failed node comprises detecting a failed node in said first group when said failed node of said first group is attempting to communicate a subfile to said at least one recipient node of said second group.

15. The method of claim 14 wherein said distribution technique adapting further comprises:

said first node communicating said subfile to said at least one recipient node of said second group.

16. The method of claim 14 wherein said distribution technique adapting further comprises:

said first node triggering a mirror node to communicate the subfile to said at least one recipient node of said second group.

17. A system comprising:

an origin node operable to partition a file F into a plurality of subfiles, wherein said plurality of subfiles correspond in number to a number of recipient nodes in a first group to which said file is to be distributed;

said origin node operable to attempt to distribute all of said plurality of subfiles to said recipient nodes, wherein said origin node attempts to distribute a different one of said plurality of subfiles to each of said recipient nodes;

said recipient nodes operable to attempt to exchange their respective subfiles received from said origin node such that each recipient node obtains all of said plurality of subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the origin node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile from the origin node;

said origin node operable to detect a failed node in said first group; and

said origin node operable to manage distribution of said file F upon detecting a failed node in said first group in a manner such that every non-failed node of said first group receives said file F .

18. The system of claim 17 wherein each of said recipient nodes are operable to attempt to distribute a subfile being received from said origin node to the others of said recipient nodes of said first group.

19. The system of claim 17 wherein said origin node is operable to attempt to distribute the plurality of subfiles to said plurality of recipient nodes of said first group concurrently.

20. The system of claim 17 wherein said said origin node is operable to trigger a mirror node to establish concurrent communication connections with non-failed nodes of said first group to communicate a subfile to said non-failed nodes.

21. A method of distributing a file from a first node to a plurality of recipient nodes, the method comprising:

attempting to distribute a plurality of subfiles that comprise a file F from a first node to a first group comprising a plurality of recipient nodes, wherein the first node attempts to distribute at least one subfile to each recipient node of said first group but not all of said plurality of subfiles are distributed from the first node to any of the recipient nodes of said first group;

said plurality of recipient nodes of said first group attempting to exchange their respective subfiles, wherein at least one recipient node of said first group begins communicating a portion of its respective subfile that it is receiving from the first node to at least one other recipient node of said first group before the at least one recipient node fully receives its respective subfile;

detecting whether one of said plurality of recipient nodes of said first group has failed; and

if a recipient node of said first group has failed, managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that the file F is distributed to each non-failed node of said plurality of recipient nodes.

22. The method of claim 21 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:

attempting to distribute a different subfile from said first node to each of said recipient nodes of said first group.

23. The method of claim 22 wherein said managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that file F is distributed to each non-failed node of said plurality of recipient nodes comprises:

said first node communicating to non-failed nodes of said first group a subfile that the first node would communicate to the failed node if the failed node were not failed.

24. The method of claim 22 wherein said managing the distribution of the plurality of subfiles to detour their distribution around the failed node such that file F is distributed to each non-failed node of said plurality of recipient nodes comprises:

said first node triggering a mirror node to communicate to non-failed nodes of said first group a subfile that the first node would communicate to the failed node if the failed node were not failed.

25. The method of claim 21 wherein said attempting to distribute the plurality of subfiles from a first node to a first group of recipient nodes comprises:

attempting to distribute the plurality of subfiles from said first node to said plurality of recipient nodes of said first group concurrently.

26. The method of claim 21 wherein said plurality of recipient nodes of said first group attempting to exchange their respective subfiles further comprises:

each of said plurality of recipient nodes attempting to establishing concurrent communication connections to every other recipient node of said first group.

27. The method of claim 21 wherein said detecting whether one of said plurality of recipient nodes of said first group has failed comprises:

said recipient nodes of said first group exchanging heartbeat messages;

at least one recipient node of said first group detecting a failed node from analysis of heartbeat messages received; and

said at least one recipient node of said first group notifying said first node of said detected failed node.

28. The method of claim 21 wherein said detecting whether one of said plurality of recipient nodes of said first group has failed comprises:

the non-failed recipient nodes of said first group sending heartbeat messages to said first node; and

said first node detecting a failed node from analysis of received heartbeat messages from the non-failed recipient nodes.